

IN THE CLAIMS:

1. (Previously Presented) A method for a storage operating system implemented in a
2 storage system to optimize an amount of readahead data retrieved from
3 a data container of the storage system, comprising:
 - 4 managing a separate set of readahead metadata for each individual read stream of
5 a plurality of read streams, wherein the read stream is a set of one or more client read
6 requests to retrieve data from a contiguous range of file offsets within a requested file;
 - 7 storing the separate set of readahead metadata in a corresponding individual
8 readset data structure for each individual read stream;
 - 9 maintaining, in the individual readset data structure, a plurality of factors for the
10 each read stream, the plurality of factors determining the amount of readahead data
11 retrieved from the data container the one or more client read requests;
 - 12 receiving a client read request associated with a particular read stream at the
13 storage system and determining an amount of readahead data to retrieve from the data
14 container for the particular read stream;
 - 15 receiving a next client read request associated with the particular read stream;
16 locating a readset data structure for the particular read stream;
17 determining whether the storage operating system is permitted to retrieve
18 readahead data from the data container in response to the received next client read
19 request; and
20 if it is determined that the storage operating system is permitted to retrieve
21 readahead data from the data container, performing the steps of,
 - 22 (i) modifying one or more of the plurality of factors stored within the readset data
23 structure for the particular read stream;
 - 24 (ii) adjusting, in response to modifying the one or more factors, an amount of
25 readahead data to retrieve for the next client read request associated with the particular
26 read stream, wherein the amount of readahead data to retrieve for the next client read

- 27 request is different from the amount of readahead data retrieved for the client read
28 request;
29 (iii) retrieving the adjusted amount of readahead data from the data container; and
30 (iv) determining if the readset data structure meets a criteria for being updated,
31 and if the readset data structure meets the criteria then updating the readset data structure.
- 1 2. (Original) The method of claim 1, wherein the data container is a file, directory, vdisk
2 or lun.
- 1 3. (Original) The method of claim 1, wherein the storage operating system is determined
2 to be permitted to retrieve readahead data from the data container when the client-
3 requested data extends the read stream past a predetermined next readahead value.
- 1 4. (Original) The method of claim 3, wherein the predetermined next readahead value is
2 stored in a readset data structure associated with the read stream.
- 1 5. (Previously Presented) The method of claim 3, wherein the predetermined next
2 readahead value is updated based on a percentage of the adjusted amount of readahead
3 data.
- 1 6. (Previously Presented) The method of claim 1, wherein a read-access style associated
2 with the data container is one of the plurality of factors used to select the amount of
3 readahead data.
- 1 7. (Previously Presented) The method of claim 6, wherein the adjusted amount of
2 readahead data equals zero if the read-access style corresponds to a random read-access
3 style.

- 1 8. (Previously Presented) The method of claim 1, wherein a number of client read
2 requests processed in the read stream is one of the plurality of factors used to select the
3 amount of readahead data.

- 1 9. (Original) The method of claim 8, wherein the number of client read requests
2 processed in the read stream is stored as a count value in a readset data structure
3 associated with the read stream.

- 1 10. (Previously Presented) The method of claim 1, wherein the amount of client-
2 requested data is one of the plurality of factors used to select the amount of readahead
3 data.

- 1 11. (Previously Presented) The method of claim 10, wherein the adjusted amount of
2 readahead data is set equal to a predetermined upper limit for large amounts of client-
3 requested data.

- 1 12. (Previously Presented) The method of claim 1, wherein the adjusted amount of
2 readahead data is doubled if the number of client read requests processed in the read
3 stream is greater than a first threshold value.

- 1 13. (Original) The method of claim 1, wherein the client-requested data is identified as
2 read-once data when either (i) the number of client read requests processed in the read
3 stream is greater than a second threshold value or (ii) a set of metadata associated with
4 the read stream indicates that the client-requested data is read-once data.

- 1 14. (Previously Presented) The method of claim 1, wherein the adjusted amount of
2 readahead data is stored in one or more buffers enqueued on a flush queue, the flush
3 queue being configured to reuse buffers after a predetermined period of time.

- 1 15. (Original) The method of claim 14, wherein the predetermined period of time equals
- 2 two seconds.

- 1 16. (Currently Amended) An apparatus configured to implement a storage operating
- 2 system that optimizes an amount of readahead data retrieved from a data container of the
- 3 apparatus, comprising:
 - 4 means for receiving a first data read command associated with a particular read
 - 5 stream, wherein the read stream is a set of one or more client read commands to retrieve
 - 6 data from a contiguous range of file offsets within a requested file;
 - 7 ;
 - 8 means for maintaining, for a selected file of a plurality of files, a plurality of
 - 9 readset data structures, each readset data structure holding a plurality of factors for a
 - 10 particular read stream, the plurality of factors allowing the system to adjust adaptively the
 - 11 amount of data retrieved from the data container;
 - 12 means for receiving a client read request associated with the particular read
 - 13 stream and means for determining an amount of readahead data to retrieve from the data
 - 14 container for the particular read stream;
 - 15 means for receiving a next client read request associated with the particular read
 - 16 stream;
 - 17 means for locating a readset data structure for the particular read stream;
 - 18 means for adjusting an amount of readahead data to retrieve from the data
 - 19 container based on modifying one or more of the plurality of factors stored within the
 - 20 readset data structure, wherein the adjusted amount is a different amount than an amount
 - 21 retrieved from the client read request;
 - 22 means for retrieving the adjusted amount of readahead data from the data
 - 23 container; and
 - 24 means for determining if the readset data structure meets a criteria for being
 - 25 updated, and if the readset data structure meets the criteria then updating the readset data
 - 26 structure.

27 means for managing a separate set of readahead metadata for each individual read
28 stream of a plurality of read streams, wherein the read stream is a set of one or more
29 client read requests to retrieve data from a contiguous range of file offsets within a
30 requested file;
31 means for storing the separate set of readahead metadata in a corresponding
32 individual readset data structure for the each individual read stream;
33 means for maintaining, in the individual readset data structure, a plurality of
34 factors for the each read stream, the plurality of factors determining the amount of
35 readahead data retrieved from the data container the one or more client read requests;
36 means for receiving a client read request associated with a particular read stream
37 at the storage system and determining an amount of readahead data to retrieve from the
38 data container for the particular read stream;
39 means for receiving a next client read request associated with the particular read
40 stream;
41 means for locating a readset data structure for the particular read stream;
42 means for determining whether the storage operating system is permitted to
43 retrieve readahead data from the data container in response to the received next client
44 read request; and
45 if it is determined that the storage operating system is permitted to retrieve
46 readahead data from the data container,
47 (i) means for modifying one or more of the plurality of factors stored within the
48 readset data structure for the particular read stream;
49 (ii) means for adjusting, in response to modifying the one or more factors, an
50 amount of readahead data to retrieve for the next client read request associated with the
51 particular read stream, wherein the amount of readahead data to retrieve for the next
52 client read request is different from the amount of readahead data retrieved for the client
53 read request;
54 (iii) means for retrieving the adjusted amount of readahead data from the data
55 container; and

56 (iv) means for determining if the readset data structure meets a criteria for being
57 updated, and if the readset data structure meets the criteria then updating the readset data
58 structure.

- 1 17. (Original) The apparatus of claim 16, wherein the data container is a file, directory,
2 vdisk or lun.
- 1 18. (Original) The apparatus of claim 16, wherein the storage operating system is
2 determined to be permitted to retrieve readahead data from the data container when the
3 client-requested data extends the read stream past a predetermined next readahead value.
- 1 19. (Previously Presented) The apparatus of claim 18, further comprising means for
2 updating the predetermined next readahead value based on a percentage of the adjusted
3 amount of readahead data.
- 1 20. (Previously Presented) The apparatus of claim 16, wherein the plurality of factors
2 used to select the amount of readahead data includes at least one of:
 - 3 (i) the amount of client-requested data,
 - 4 (ii) a number of client read requests processed in the read stream, and
 - 5 (iii) a read-access style associated with the data container.
- 1 21. (Previously Presented) The apparatus of claim 16, wherein the adjusted amount of
2 readahead data is doubled if the number of client read requests processed in the read
3 stream is greater than a first threshold value.
- 1 22. (Currently Amended) A storage system configured to optimize an amount of
2 readahead data retrieved from a data container of the storage system, the storage system
3 comprising:
 - 4 a network adapter for receiving a client read request for a particular read stream;

5 a memory configured to store instructions for implementing a storage operating
6 system that performs the steps of:

7 locating a readset data structure for the particular read stream;
8 adjusting, in response to the readset data structure for the particular read stream,
9 the adjusted amount of readahead data from the data container based on a plurality of
10 factors, the plurality of factors allowing the system to adjust adaptively the amount of
11 data retrieved from the data container;

12 retrieving the adaptively adjusted amount of readahead data from the data
13 container;

14 managing a separate set of readahead metadata for each individual read stream of
15 a plurality of read streams, wherein the read stream is a set of one or more client read
16 requests to retrieve data from a contiguous range of file offsets within a requested file;

17 storing the separate set of readahead metadata in a corresponding individual
18 readset data structure for the each individual read stream;

19 maintaining, in the individual readset data structure, a plurality of factors for the
20 each read stream, the plurality of factors determining the amount of readahead data
21 retrieved from the data container the one or more client read requests;

22 receiving a client read request associated with a particular read stream at the
23 storage system and determining an amount of readahead data to retrieve from the data
24 container for the particular read stream;

25 receiving a next client read request associated with the particular read stream;
26 locating a readset data structure for the particular read stream;

27 determining whether the storage operating system is permitted to retrieve
28 readahead data from the data container in response to the received next client read
29 request; and

30 if it is determined that the storage operating system is permitted to retrieve
31 readahead data from the data container, performing the steps of,

32 (i) modifying one or more of the plurality of factors stored within the readset data
33 structure for the particular read stream;

34 (ii) adjusting, in response to modifying the one or more factors, an amount of
35 readahead data to retrieve for the next client read request associated with the particular
36 read stream, wherein the amount of readahead data to retrieve for the next client read
37 request is different from the amount of readahead data retrieved for the client read
38 request;
39 (iii) retrieving the adjusted amount of readahead data from the data container; and
40 (iv) determining if the readset data structure meets a criteria for being updated,
41 and if the readset data structure meets the criteria then updating the readset data structure.

- 1 23. (Original) The storage system of claim 22, wherein the data container is a file,
2 directory, vdisk or lun.
- 1 24. (Original) The storage system of claim 22, wherein the storage operating system is
2 determined to be permitted to retrieve readahead data from the data container when the
3 client-requested data extends the read stream past a predetermined next readahead value.
- 1 25. (Previously Presented) The storage system of claim 24, wherein the predetermined
2 next readahead value is updated based on a percentage of the adjusted amount of
3 readahead data.
- 1 26. (Previously Presented) The storage system of claim 22, wherein the plurality of
2 factors used to select the amount of readahead data includes at least one of:
 - 3 (i) the amount of client-requested data,
 - 4 (ii) a number of client read requests processed in the read stream, and
 - 5 (iii) a read-access style associated with the data container.
- 1 27. (Previously Presented) The storage system of claim 22, wherein the adjusted amount
2 of readahead data is doubled if the number of client read requests processed in the read
3 stream is greater than a first threshold value.

1 28. (Currently Amended) A computer-readable media containing program instructions
2 executed by a processor, comprising:

3 program instructions that receive a first data read command associated with a
4 particular read stream, wherein the read stream is a set of one or more client read
5 commands to retrieve data from a contiguous range of file offsets within a requested file;

6 program instructions that maintain, for a selected file of a plurality of files, a
7 plurality of readset data structures, each readset data structure holding a plurality of
8 factors for a particular read stream based on a plurality of factors, the plurality of factors
9 allowing the system to adjust adaptively the amount of data retrieved from the data
10 container;

11 program instructions that receive a client read request associated with the
12 particular read stream at the storage system and determining an amount of readahead data
13 to retrieve from the data container for the particular read stream;

14 program instructions that receive a next client read request associated with the
15 particular read stream;

16 program instructions that locate a readset data structure for the particular read
17 stream;

18 program instructions that determine whether the storage operating system is
19 permitted to retrieve readahead data from the data container in response to the received
20 next client read request;

21 if it is determined that the storage operating system is permitted to retrieve
22 readahead data from the data container, executing program instructions that
23 (i) adjust an amount of readahead data to retrieve from the data container based on
24 modifying one or more of the plurality of factors stored within the readset data structure;
25 (ii) retrieve the adjusted amount of readahead data from the data container; and
26 (iii) determine if the readset data structure meets a criteria for being updated, and
27 if the readset data structure meets the criteria then updating the readset data structure.

28 program instructions that manage a separate set of readahead metadata for each
29 individual read stream of a plurality of read streams, wherein the read stream is a set of

30 one or more client read requests to retrieve data from a contiguous range of file offsets
31 within a requested file;
32 program instructions that store the separate set of readahead metadata in a
33 corresponding individual readset data structure for the each individual read stream;
34 program instructions that maintain, in the individual readset data structure, a
35 plurality of factors for the each read stream, the plurality of factors determining the
36 amount of readahead data retrieved from the data container the one or more client read
37 requests;
38 program instructions that receive a client read request associated with a particular
39 read stream at the storage system and determining an amount of readahead data to
40 retrieve from the data container for the particular read stream;
41 program instructions that receive a next client read request associated with the
42 particular read stream;
43 program instructions that locate a readset data structure for the particular read
44 stream;
45 program instructions that determine whether the storage operating system is
46 permitted to retrieve readahead data from the data container in response to the received
47 next client read request; and
48 if it is determined that the storage operating system is permitted to retrieve
49 readahead data from the data container, program instructions that,
50 (i) modify one or more of the plurality of factors stored within the readset data
51 structure for the particular read stream;
52 (ii) adjust, in response to modifying the one or more factors, an amount of
53 readahead data to retrieve for the next client read request associated with the particular
54 read stream, wherein the amount of readahead data to retrieve for the next client read
55 request is different from the amount of readahead data retrieved for the client read
56 request;
57 (iii) retrieve the adjusted amount of readahead data from the data container; and
58 (iv) determine if the readset data structure meets a criteria for being updated, and
59 if the readset data structure meets the criteria then update the readset data structure.

- 1 29. (Original) The computer-readable media of claim 28, wherein the data container is a
2 file, directory, vdisk or lun.
- 1 30. (Previously Presented) The method of claim 1, wherein the retrieved readahead data
2 is stored in one or more buffers, the buffers containing a flush queue, the flush queue
3 being configured to reuse buffers after a predetermined period of time.
- 1 31. (Previously Presented) The method of claim 30, wherein the read stream
2 corresponds to a read-once data transfer and data retrieved from the data container is
3 stored in the flush queue.
- 1 32. (Previously Presented) The method of claim 30, wherein the retrieved readahead
2 data is stored in the flush queue.
- 1 33. (Previously Presented) The method of claim 30, wherein one or more buffers
2 accessed from the flush queue are re-enqueued on a normal queue.
- 1 34. (Cancelled)
- 1 35. (Cancelled)
- 1 36. (Cancelled)
- 1 37. (Cancelled)
- 1 38. (Cancelled)
- 1 39. (Cancelled)

1 40. (Cancelled)

1 41. (Cancelled)

1 42. (Cancelled)

1 43. (Cancelled)

1 44. (Cancelled)

1 45. (Cancelled)

1 46. (Cancelled)

1 47. (Cancelled)

1 48. (Cancelled)

1 49. (Cancelled)

1 50. (Cancelled)

1 51. (Cancelled)

1 52. (Cancelled)

1 53. (Cancelled)

1 54. (Cancelled)

1 55. (Cancelled)

1 56. (Cancelled)

1 57. (Cancelled)

1 58. (Cancelled)

1 59. (Cancelled)

1 60. (Cancelled)

1 61. (Cancelled)

1 62. (Cancelled)

1 63. (Cancelled)

1 64. (Cancelled)

1 65. (Cancelled)

1 66. (Cancelled)

1 67. (Cancelled)

1 68. (Cancelled)

1 69. (Cancelled)

1 70. (Cancelled)

1 71. (Cancelled)

1 72. (Cancelled)

1 73. (Cancelled)

1 74. (Cancelled)

1 75. (Cancelled)